Enable Portrait Privacy Protection in Photo Capturing and Sharing

Lan Zhang*, Kebin Liu *, Xiang-Yang Li†, Puchun Feng*, Cihang Liu*, Yunhao Liu*
* School of Software, Tsinghua University
† Department of Computer Science, Illinois Institute of Technology

arXiv:1410.6582v1 [cs.CR] 24 Oct 2014

***Abstract*—The wide adoption of wearable smart devices with onboard cameras greatly increases people's concern on privacy infringement. Here we explore the possibility of easing persons from photos captured by smart devices according to their privacy protection requirements. To make this work, we need to address two challenges: 1) how to let users explicitly express their privacy protection intention, and 2) how to associate the privacy requirements with persons in captured photos accurately and efficiently. Furthermore, the association process itself should not cause portrait information leakage and should be accomplished in a privacy-preserving way. In this work, we design, develop, and evaluate a protocol, called *InvisibleMe*, that enables a user to flexibly express her privacy requirement and empowers the photo service provider (or image taker) to exert the privacy protection policy. Leveraging the visual distinguishability of people in the field-of-view and the dimension-order-independent property of vector similarity measurement, InvisibleMe achieves high accuracy and low overhead. We implement a prototype system, and our evaluation results on both the trace-driven and real-life experiments confirm the feasibility and efficiency of our system.**

I. Introduction

Nowadays, smart devices with onboard cameras *e.g.*, smart phones and glasses, are pervasive in our daily lives. These smart devices can capture and even share photos without informing the parties in the picture, thus raises many concerns on people's privacy infringement. Particularly, the wilder adoption of smart glasses, *e.g.* Google Glass, leads to severe concerns for misuse because Glass can capture photos/videos far less conspicuously than a traditional hand-held device. Secretive photographing without clear warning beforehand and possession of secretively taken photos are both privacy violations. Even worse, if the photos which contain information beyond what users want to reveal are shared in Internet, it will make users extremely susceptible to various attacks.

To protect people's portrait privacy from unwilling photo-taking and publication, many photo service providers or users have taken actions in different ways. For example, some Glass wearers whip their device off in inappropriate situations, such as in gym locker rooms or work meetings; some business bans smart glasses inside their buildings to respect customers' privacy [1]; and the Glass manufacturer (*e.g.,* Google) does not allow developers to create applications that take photo silently. Lawmakers are also beginning to consider various privacy issues of Glass, including whether it should be capable of facial recognition [2]. Although face recognition is a useful function, especially for social applications, face is a critical private identifiable information. At present, there are no facial recognition technologies built into Glass and the manufacturer has no plans to use it unless they have strong privacy protections in place. These methods, however, are broad-brush and blunt which can significantly hurt the applications of smart devices. Therefore, it is appealing to consider how one might build a system in which users do not leak portrait privacy while guaranteeing a comfortable usage of smart glasses/cameras.

Instead of discarding the smart glasses/cameras due to privacy concerns, in this work, we seek a solution for reaching an ultimate goal of privacy-friendly Glass/camera, operating transparently to end users. Our solution will let end users to express their privacy requirements and glasses/cameras or photo service providers will exert the privacy protection mechanisms. When taking a photo/video, the smart device will detect who (in the picture/video) requested privacy protection, and then remove them from the image automatically. Our protocol can also be used for automatically tagging people in a photo when a user expresses an interest to be tagged with his/her information.

To implement such a privacy-friendly camera, we need to address several critical challenges. **First**, we should enable privacy-advocator efficiently and flexibly to express his/her privacy requirements/intentions. Several methods could be used for this task, such as using visible specialized tag (*e.g.*, QR code), or encoding the request and transmitting it using wireless devices. For users' convenience and aesthetics, in this work, we adopt the latter approach by encoding his/her portrait in the request. Then, the **second** challenge is that we should accurately and efficiently associate each privacy-seeking user with an image region in the photo taken by another user (the photographer). Furthermore, the association process itself should not cause portrait information leakage and should be accomplished in a privacy-preserving way. Face recognition [3], [4] is widely used to identify people in photos, but in practice it suffers when there lacks a clear front view of faces. Sophisticated but complicated matching schemes may cause high overhead and long delay. The matching problem itself is difficult due to the accuracy and efficiency requirements, let alone completing matching process in a private and non-interactive manner with untrusted server. Matching a user's privacy-expression with a possible people in a photo can be reduced to some sort of vector matching. Many private vector matching protocols use multi-party computation techniques, which require frequent interactions among participants. Most

exiting private vector matching methods (in both multi-party computation and outsourced manner) use homomorphic encryption [5], [6] or garble circuit [6], and cause high computation cost for both client and cloud. The **third** challenge is that the privacy-friendly Glass/camera should be transparent to all users and cause minimal extra overhead to mobile devices. An ad hoc approach may lead to requirements for "always-on" neighbour discovery, frequent information exchanging as well as complex image matching computation on user devices. To reduce the overhead of users, our protocol will outsource most of these tasks to cloud with a well-designed strategy to prevent privacy leakage to untrusted cloud and other users.

The main contributions of this work are as follows:

•To the best of our knowledge, we are the first to present a portrait privacy preserving photo capturing and sharing approach. With our approach in Section III, people who require not to be captured in photo will be automatically erased from the photo and verification of the removal is also supported.

•We comprehensively analyze the privacy issues during the photo capturing and sharing and define three types of threats in Section II. Based on the proposed model, we present a solution protecting all three types of privacy information.

•For accurate and efficient matching between people's privacy intentions and people in the photo, we introduce a graph-based portrait profile and design a robust matching algorithm to recognize and erase privacy-seeking people in Section IV.

•We propose a highly efficient privacy-preserving vector distance protocol in a non-interactive manner with untrusted server in Section V, which significantly reduces the computation complexity and communication cost than existing homomorphic encryption and garble circuit based solutions. With our protocol, most computation tasks are transferred from smart devices to the cloud in a privacy-preserving way.

•We design and implement a prototype system and verify the effectiveness of our scheme by extensive experiments as well as case studies in Section VI.

II. Privacy Requirement

A. Motivation

To protect users' portrait privacy, one straight-forward approach that has already been adopted by business and manufacturer is to simply suppress the usage of smart cameras in specific place and time, *e.g.*, turn off the photographing functionality in a meeting room or forbid silent photo taking. These methods, however, are broad-brush and blunt which can significantly hurt the applications of smart devices. Besides, they do not meet the user varying privacy intentions. In fact, not all persons are unwilling to be photographed and also people will feel quite uncomfortable if they are frequently forced to turn off their Glass. Thus, in this work we seek solution to give privacy control back to persons being photographed in a smarter way.

As an example shown in Fig. 1, when someone uses his smart Glass to take a photo, people in the field of view (FOV) should be notified (or the photographer should know the privacy protection intentions of people in FOV). Then persons who are unwilling to be photographed, *e.g.* Neighbor 1, should

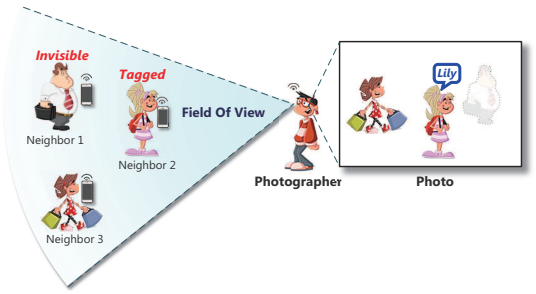

Fig. 1: Example application scenario of *InvisibleMe*: the invisible person is erased from the photo and the tagged person is labeled in the photo.

have a convenient way to specify their privacy intentions, and thus be automatically erased from the photo. We refer to them as *invisible users*. Users that would like to make friends with the photographer, *e.g.*, Neighbor 2, can be automatically tagged on the photo and share information. We refer to them as *tagged users*. The photographer just takes other people into the photo as usual. The system can motivate the photographer by mutual registry: only those who protect other invisible users can be registered to be protected by others. Supporting tagging people automatically, which could be helpful and fun in many scenarios (*e.g.*, facebook), also gives incentives to the photographer. Besides, after one time setting the solution should be transparent to all users and avoid incurring high overhead to their smart devices. Finally, the portrait privacy protection strategy should be well designed and avoid further leakage of any type of private information.

B. Threats to Portrait Privacy

Photos contain rich information, including people's appearance, location, activities, *etc.*. Facing massive cameras and image analysis techniques [4], [7] , people's portrait privacy is badly in need of protection. In this work, we focus on protecting users' portrait information. Here, a portrait not only includes the user's face but also his/her body, since clothes and accessories could also reveal identification information. We consider three types of threats to portrait privacy.

Visual portrait privacy. The most intuitive way to violate a user's portrait privacy is to capture and publish (*e.g.*, through photo sharing systems) a photo containing his/her visible portrait. Simply blurring all faces in images, *e.g.*, [8] and Google Street View, will disable the normal photographing function.

In our protocol InvisibleMe we propose to match the people in the photo to their privacy protection intentions, and erase only people that should be invisible. As we will discuss in detail in Section III, a user express his/her privacy requirement by encoding his/her portrait, which clearly cannot be transmitted in its original form (otherwise his/her portrait privacy is broken by himself/herself). So, we need to provide privacy protections in all these operations.

Portrait feature privacy. This type of threats occur inside some image services, *e.g.* image matching or face recognition. These services don't use visible images of directly, but take

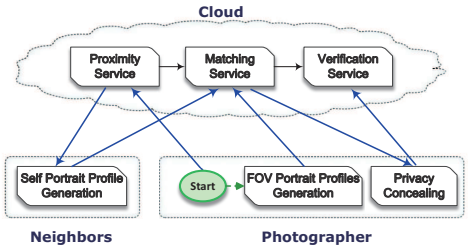

Fig. 2: Baseline system architecture.

feature vectors of image as the descriptor, *e.g.*, Eigenfaces [4] and color histogram. But users can also be identified by features of portrait image. For example, face images can be reconstructed from face vectors [9]. During the process, the leakage of portrait features also violates users' privacy.

Inference privacy. Even if an image system hides original images and other information such as their feature vectors, an adversary with a collection of images (an image dictionary) can infer the hidden content using the similarity measurement function of the system. Hence, we should prevent adversaries from obtaining the similarity measuring results to enhance the privacy protection.

InvisibleMe leverages existing solutions to protect other user privacy, *e.g.*, location privacy [10], since it is not the focus of this study.

Adversary model. Our approach defends a user's portrait privacy against both untrusted cloud server and malicious users. For the cloud, we apply the widely used "honest-but-curious" assumption. The cloud server will follow the protocol, but might conduct extra work to harvest portrait images of invisible people, reconstruct invisible portraits using feature vectors or infer the invisible content using image dictionary. Also, we assume the cloud won't collude with any client to conduct an attack. A malicious user could participate to harvest other users' portrait information by eavesdropping their communication with the cloud. All users except the photographer should be prevented from obtaining the portrait information of invisible users. Although the photographer already owns the photo of invisible people, he/she may misbehave to preserve the invisible people who should be erased and publish the photo through Internet. So, we also need a verification scheme against dishonest photographers.

III. System Design Overview

Facing critical challenges introduced in Section I, we design our system to achieve both the privacy and system efficiency goals. With our graph-based portrait matching algorithm, the baseline approach is effective to protect users' visual portrait privacy by accurately locating invisible people in photos and erases them automatically. But there is a risk of exposing portrait features to untrusted cloud and other participants. Furthermore, we propose an efficient privacy-preserving outsourced vector distance protocol, based on which an advanced approach provides portrait feature privacy and inference privacy protection with little extra overhead for the client. In this section, we will sketch our system architecture.

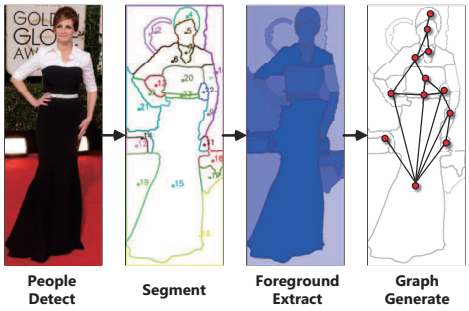

Fig. 3: Portrait graph representation.

A. Overview of Baseline System

In our system, there are three parties involved: the photographer, who takes the photo; the neighbors, who could be in the FOV of the photographer (as presented in Fig. 1); the cloud, who takes charge of location, communication and computation services. The architecture and workflow of our baseline system are illustrated in Fig. 2

We would like to take a typical photographing process as an example to describe the functionality of each component and the system workflow. Users create their personal portrait profiles using the *Self Portrait Profile Generation* component and encode their portrait profiles to express their privacy requirements. In our design, a set of vision feature vectors are extracted from subregions of a portrait image. Both face and body features are extracted, in case that there may lack a clear front view of face. We introduce a graph structure to encode the extracted vision features (feature vectors are properties of nodes) and use the graph as the portrait profile, as the examples in Fig. 3. Compared with uploading the original image, the feature graph shows a low risk on privacy leakage without lose of matching functionality, and are much more efficient for both computation and communication. Besides, the graph representation is highly robust for pose changes of people and cameras. For each invisible user, his/her self portrait profile is generated once and for all until he/she updates it. While the face features of a user remains the same, the user could change his/her outfits. The portrait profile could be automatically updated when the user selfies or while he/she uses the phone with the frontal camera facing himself/herself. In our advanced approach, transformed version of portrait graph are used to improve privacy protection. With portrait graphs, we convert the people matching problem to graph matching problem.

Triggered by a photo shooting action, the *Proximity Service* on cloud will be automatically notified and start to check if there are invisible people in the FOV of the photographer. If any, the cloud will inform them and start the next step. Proximity service can be realized easily using common location service and onboard compass. By restricting the number of potential matched invisible users, our system can achieve high matching accuracy and low overhead.

In the next step, after being informed by cloud, invisible neighbors upload their self portrait profiles to cloud (this could be done in advance to reduce the delay). Meanwhile, the photographer detects all people in the photo and extracts their

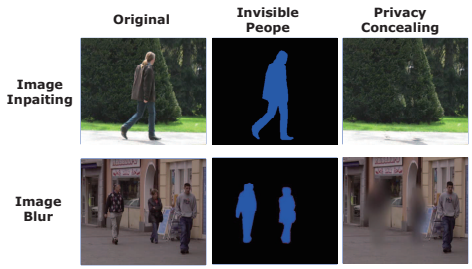

Fig. 4: Example of automatic privacy concealing (erase invisible people from photos) by image inpainting [12] or blur.

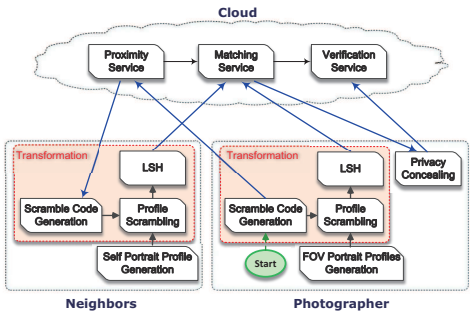

Fig. 5: Advanced system architecture.

portrait profiles with the *FOV Portrait Profiles Generation* component, which works similarly to self portrait profile generation. These profiles will be uploaded to cloud as well. Then the *Matching Service* will match portrait profiles of invisible users to portrait profiles from the photo and determine people that should be erased from the photo. The graph matching algorithm will be discussed in detail in Section IV. The matched results will be sent to photographer, and then the *Privacy Concealing* component will erase the corresponding image regions of invisible people from the photo automatically by blurring or other more sophisticated techniques, like image inpainting [11], [12], to maximize the aesthetics. We show an example of removing invisible people from photo in Fig. 4. For the inpainting, we use the code from Criminisi's work [12]. After the removal, the photographer can store or share the photo using the cloud service. Note that, based on the personal specification, the whole procedure works the same way for tagged users and the "erase" operation can be alternated to "tag" to augment many social applications.

In case there are dishonest photographers who don't complete the removal, InvisibleMe supports verification of removal. All invisible users' portrait profiles have been uploaded to the cloud in the previous stage. Once the photo is shared through Internet, the *Verification Service* will check the photo as follows the cloud first conducts a people detection on the photo and extracts all portrait profiles; then the cloud matches these profiles with the cached profiles of invisible neighbors, if there is a matching, it can tell that the photographer didn't follow the protocol. The verification process can be completed alone by the cloud without any interaction with users.

B. Overview of Advanced System

The baseline protocol protects the visual privacy of people's portraits, but exposes users' portrait profile (*i.e.* feature vectors) to the cloud and even the eavesdroppers. With some feature vectors an adversary could have a chance to match them with existing photos or even reconstruct the photo. In the advance approach, we retains the visual portrait privacy protection and improve the system to protect users' portrait feature privacy and inference privacy (defined in Section II-B). The graph based profile matching scheme should be conducted in a privacy-preserving manner. The core of the portrait profile matching algorithm is to measure the distance between vision feature vectors. We cannot directly adopt existing privacy-preserving vector distance protocols based on homomorphic encryption [5], [6] and gabled circuit [6] due to their large computation and communication cost. In InvisibleMe we propose a highly efficient outsourced vector distance protocol (see Section V). As shown in the red blocks in Fig. 5, combining a well designed scrambling scheme and locality sensitive hash, all invisible neighbors can secretly transform their vectors in a distance preserving way. Then the cloud can measure vector distances using the transformed vectors and match transformed portrait graph with the same algorithm as in the baseline system. Our scheme protects invisible users' portrait profiles (from both himself/herself and the photographer) with very little extra cost on the client side for generating transformed portrait graph. But, it saves computation cost for the cloud, because the distance computation of high-dimensional real number vectors is converted to distance of low-dimensional binary hash code.

The architecture of the advance system is presented in Figure 5, except vector transformation and verification, other components are the same as that in the baseline system. Here, the verification is more challenging, because the cloud only knows the transformed portrait graphs of invisible users. Without knowing the secret transformation, the cloud cannot compare them with portrait graphs directly extracted from the uploaded photo. When an invisible user needs to check if his/her portrait has been removed, he/she need to start a verification and participate in the process as follows: the cloud sends all extracted feature vectors in a random order to the invisible user. Note that, these feature vectors are supposed to belong to preserved visible people if the photographer is honest. And the invisible user transforms them in the same way as his/her self feature vectors and sends the results transformed vectors to the cloud. Then the cloud can compare preserved people in the photo and invisible neighbors using transformed portrait graphs, and detect the dishonest photographer.

IV. Portrait Profile Generation and Matching

A. Portrait Profile Generation

After applying a people detection on a photo [7], [13], we obtain portrait images (including both people faces and bodies) from this photo as shown in the first subfigure of Fig. 3. Then portrait image can be segmented into adjacent regions by different colors and textures [14]. Given one portrait image, a graph $G = (V, E)$ can be constructed, where V is a set of nodes representing segmented regions and E are edges connecting any two regions that share a boundary. Then

we measure each node's confidence of being a part of the person and remove the node with low confidence to eliminate the background. The confidence calculation is omitted due to space limitation. Fig. 3 shows examples of foreground extraction and portrait graph generation, which provide more accurate graph representation of people portrait. The result of foreground extraction can also be employed by the privacy concealing component as the accurate erase area to achieve better looking removal, as shown in Fig. 4. For each region of portrait, vision feature vectors, *e.g.*, face feature vector, color histogram and texture vector, are extracted as property of the corresponding node. We will give more detail about node properties in the implementation section (Section. VI).

B. Portrait Graph Matching Scheme

To achieve accurate and efficient portrait graph matching, there are several challenging issues should be addressed with low computation cost: graphs structure of the same person varies due to changing illumination condition and viewpoint; incomplete graphs could be produced due to occlusion; portrait profile could still contain some noisy nodes from background. As a result, the matching algorithm should be elastic to node/edge division, aggregation, insertion and deletion, and robust to noise nodes. Existing graph matching methods usually have application-oriented specifications [15]–[17], *e.g.*, assumptions about node numbers, graph structure and pre-knowledge of correspondences, make them difficult to be directly applied in this work. To meet the critical requirements of portrait profile matching, we design a voting based strategy in which both the node similarity and graph structure are considered.

Let graph $G^x = (V^x, E^x)$ denote portrait profile X (say produced by a user) and $G^y = (V^y, E^y)$ denote portrait profile Y (say produced by a photographer). Here $V^x = \{v_1^x, v_2^x, ..., v_p^x\}$ and $V^y = \{v_1^y, v_2^y, ..., v_q^y\}$. Each node own some feature vectors as its property. In order to improve matching accuracy as well as speed up the computation process, we add a *label* for each node, which describes its type, for example, *human face* or *human body*. Only nodes of the same type can be matched. The distance between two nodes of different types is regarded as infinite. As human face is a strong feature to identify a person, our matching scheme will firstly consider the matching between nodes labeled with "human face" (*e.g.*, Node No.5 in Fig. 3), then invoke an integrated graph matching. In this way, our method provides more accurate and robust matching than existing face recognition based methods.

Initialization. Let the similarity between nodes v_i^x and v_j^y be $\mathbf{S}(v_i^x, v_i^y)$, which can be obtained through measuring the distances between feature vectors of two graph nodes. Note that, if v_i^x and v_i^y have different type labels, $\mathbf{S}(v_i^x, v_i^y)$ is set to zero. In Section V, we will discuss the details of privacy preserving vector distance computation. During the matching process, a matrix M with p rows and q columns is built. Each entry $M_{ij} = \{f_{ij}, n_{ij}, c_{ij}\}$ of the matrix is a triple where f_{ij} is a boolean flag indicating whether node v_j^y is a possible match for node v_i^x, n_{ij} caches the one-hop neighbor match information and c_{ij} is a counter. Details of these parameters will be presented in the following parts. A match is represented as an assignment for all $\{f_{ij}\}$, where there is at most one f_{ij} equaling $TRUE$ for every column j. All $\{f_{ij}\}$ are initiated to $TRUE$.

After the initialization, our graph matching scheme consists of three stages.

Stage 1. We eliminate wrong matches based on the similarities of node pairs. If $\mathbf{S}(v_i^x, v_i^y)$ is above a pre-specified threshold ξ_s, the corresponding flag f_{ij} is set to $TRUE$, otherwise, we eliminate this match by set f_{ij} to $FALSE$. Note that, a node in V^x does not necessarily have a possible match in V^y, thus there can be rows with all $FALSE$ flags. After this stage, all node pairs with $TRUE$ flags are considered as *candidate matches*.

Stage 2. We explore the one-hop neighbor matching for each candidate match. For each candidate match (v_i^x, v_j^y), the neighbor sets of them are denoted as $NE(v_i^x)$ and $NE(v_j^y)$. We find the most likely mapping from $NE(v_i^x)$ to $NE(v_j^y)$. To achieve this, we firstly look for potential matches in matrix M for each node in $NE(v_i^x)$. We then connect each node in $NE(v_i^x)$ with its matched nodes in $NE(v_j^y)$ with undirected edges. Nodes in both sets as well as the edges form a bipartite graph and the problem can be transformed to find a maximum match on the bipartite graph. To address this problem, we apply the Hungary algorithm [18] which outputs a mapping from $NE(v_i^x)$ to $NE(v_j^y)$. As mentioned above, the mapping is denoted as n_{ij}.

$$n_{ij}(v_a^x) = \begin{cases} v_b^y & \text{if } v_a^x \text{ matches } v_b^y \\ \Phi & \text{if there is no match in } NE(v_j^y) \text{ for } v_a^x \end{cases}$$

where $v_a^x \in NE(v_i^x)$ and $v_b^y \in NE(v_j^y)$.

Stage 3. We choose at most one assignment for each node in V^x by a voting based scheme. For each candidate match (v_i^x, v_j^y), we build two trees rooted at v_i^x and v_j^y on graph G_x and G_y respectively. The two trees are traced in parallel on two graphs with the BFS method. Here we restrict the tree growth to the constraint that, once a node v_k^x on G_x and its matched node v_g^y are appended to the trees, the neighbors of v_k^x which have not been included can be added to the tree only if they have matched nodes in $NE(v_g^y)$ according the recorded mapping n_{kg}. When two trees have grown to the maximum size, we get a possible match for the subgraphs. In this approach, we propose a voting scheme to determine the best match. That is, for each candidate match (v_k^x, v_g^y) on the two trees, we increase the counter value of c_{kg} in entry M_{kg}. After trees of all candidate matches (v_i^x, v_j^y) voted, we check the c_{ij} in each entry M_{ij} and retain the largest one for each column. Then matrix M indicates a most likely match of G^x and G^y and the similarity between the two portrait profiles are calculated by integrating similarities of all matched nodes and edges.

$$\mathbf{S}(G^x, G^y) = \frac{\sum_{f_{ij}=\text{TRUE}} \mathbf{S}(v_i^x, v_j^y)}{\parallel V^x \parallel + \parallel V^y \parallel} + \frac{\sum_{e_{ab} \in E^x} \sum_{e_{cd} \in E^y} \delta(e_{ab}, e_{cd})}{\parallel E^x \parallel + \parallel E^y \parallel}$$

where

$$\delta(e_{ab}, e_{cd}) = \begin{cases} 1 & \text{if } f_{ac} = TRUE \ \& \ f_{bd} = TRUE \\ 0 & \text{otherwise} \end{cases}$$

V. Improvement of Preserving Portrait Privacy

The main idea of our outsourced privacy-preserving vector distance protocol is to transform the original vectors to random vectors, meanwhile, preserve the distance among vectors. Moreover, the transformation should be kept secret from adversaries. In this way, the distance can be measured on transformed vectors as on original vectors (which means light-weight computation), but the adversary cannot obtain the original vectors nor compute the distance between the transformed vectors and vectors from a dictionary to infer the original ones. Let the distance function of two vectors $\mathbf{x} = (\mathbf{x}_1, \mathbf{x}_2, \cdots), \mathbf{y} = (\mathbf{y}_1, \mathbf{y}_2, \cdots) \in \mathbf{V}$ be $\mathsf{d}(\mathbf{x}, \mathbf{y})$. As shown in Fig. 5, we design the transformation with two main building blocks: *Profile Scrambling* and *Locality Sensitive Hash (LSH)*. The profile scrambling module works based on the observation that vector distances are *dimension-order-independent*, that is when we randomly change the dimension order of both $\mathbf{x}$ and $\mathbf{y}$ consistently to obtain scrambled $\mathbf{x}'$ and $\mathbf{y}'$, we have $\mathsf{d}(\mathbf{x}, \mathbf{y}) \equiv \mathsf{d}(\mathbf{x}', \mathbf{y}')$. Once the scrambling order is kept secret, the original vectors are protected and a dictionary based inference is prevented. In case there may be some dimension-dependent characteristics of vision feature vectors, *e.g.*, in the color histogram the dimensions representing red component usually have large values, we employ the LSH module to transform the scrambled feature vectors into another low-dimensional vector space. LSH hides the scrambled feature vectors from all parties and makes the statistic analysis on scrambled vectors infeasible, meanwhile it also preserves the distance among vectors. Besides, lower-dimension vectors reduce the cost for vector distance computation and vector transmission. On the other hand, changing the dimension order of $\mathbf{x}$ randomly to $\mathbf{x}'$ makes their hashes totally different, because there is a random distance between them. Hence, the vector scrambling works like a random salt to strength the security property of LSH as well, that makes the dictionary attack against LSH infeasible. Combining vector scrambling and LSH, we protect feature vectors of invisible users from untrusted cloud and other parties, and outsource most computation to the cloud in a secure and noninteractive manner.

In the following subsections, we will introduce the LSH based vector distance measurement as a preliminary, then present our privacy-preserving vector distance protocol.

A. LSH based Vector Distance Measurement

The key insight behind LSH is that it is possible to construct hash functions such that close vectors will have the same hash value with higher probability than vectors that are far apart. Different LSH functions are designed for various distance metrics, *e.g.*, Euclidean distance, Hamming distance, cosine distance. In our system, we use the commonly used Euclidean distance $\mathsf{d}_E(\mathbf{x}, \mathbf{y}) = \sqrt{\sum_i (\mathbf{x}_i - \mathbf{y}_i)^2}$. Particularly, for high-dimensional vector space $\mathbf{V} = R^D$ with Euclidean distance, an LSH function is defined as follows [19]:

$$\mathbf{H}(\mathbf{x}) = < h_1(\mathbf{x}), h_2(\mathbf{x}), \cdots, h_m(\mathbf{x}) > \quad (1)$$

$$h_i(\mathbf{x}) = \begin{cases} 1 & \text{if } \frac{\mathbf{a} \cdot \mathbf{x}}{W} \geq 1, i = 1, 2, \cdots, m \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

where $\mathbf{a} \in R^D$ is a random vector with each dimension chosen independently from the standard Gaussian distribution $N(0, 1)$. Here each h_i is an atomic LSH function, and the LSH function $\mathbf{H}$ generates a hash vector of the input vector by concatenating m scalar atomic hash values. The window size W and m control the distance range that the mapping is sensitive to. In the advance system, the cloud determines the hashing function $\mathbf{H}$ and publishes it to all participants.

According to the definition of LSH, the differences between hash vectors indicate the distance between original vectors. In this work we apply the Hamming distance between hash vectors to approximate the distance between original vectors. Our experimental results show that the Hamming distance between hash vectors is nearly monotonic to distance measurement between original vectors.

B. Outsourced Privacy-preserving Distance Computing

Here, we assume that each user share a secure communication channel with the cloud. Then the transformed vectors are protected from other participants. In a specific round of photographing, to preserve distance between transformed vectors, the challenge is that all participants (photographer and invisible neighbors) must scramble their feature vectors in a consistent order individually and secretly. We refer to the scramble order as scramble code $\mathbf{SC}$. To achieve the same $\mathbf{SC}$, all participants first need to generate a same random seed R secretly. The multi-user agreement protocol requires that the untrusted cloud cannot learn the random seed and the scramble code, although it controls all communications between users.

Random number exchange. There are many well-designed group key agreement protocols [20], [21], but most of them require multiple communication rounds among participants, which could cause long delay. Utilizing the honest-but-curious cloud and secure communication channels between the cloud and users, we adapt the practical distributed group key agreement protocol proposed in [20] to achieve round optimum and computation efficient random number agreement. Let $U_1, \cdots, U_n$ be a dynamic subset of all users who want to generate a common random number and our protocol is presented in Algorithm 1. With this protocol, the photographer and his/her invisible neighbors can obtain the same random number, while the cloud learns nothing about the random number.

Scramble code generation. After obtaining consistent random seed R, each participant generates the scramble code using Algorithm 2, and rearranges the dimension order of each feature vector according to the scramble code.

As illustrated in Fig. 5, after scrambling feature vectors, the photographer and invisible neighbors apply LSH to scrambled vectors to get transformed vectors for current round. The cloud can simply use the transformed vectors to compute distance and conduct the same graph matching algorithm as in the basic scheme. While the membership doesn't change, the random number remains the same. In this case, the photographer can use the same random number to generate transformed vectors for new photos, and all invisible neighbors do not need any

Algorithm 1 Random Number Agreement.

System Initialization:

Cloud generates and publishes system parameters: 1) a large prime number $p = \Theta(2^{cN})$, a constant $c \geq 1$, $q = \Theta(2^N)$ and $g \in Z_p$ of order $q = \Theta(2^N)$.

Each user U_i generates his private parameter $a_i \in Z_q$ and public parameter $b_i = g^{a_i} \mod p$ and sends b_i to the cloud.

Cloud checks that $b_i^q \equiv 1 \mod p$ for all $i = 1, \cdots, n$.

Runtime:

1: Cloud arranges n users' indices in a cycle and sends b_{i-1} and b_{i+1} to each $U_i, i = 1, \cdots, n$.

2: Each $U_i, i = 1, \cdots, n$ computes c_i and sends it to cloud

$$c_i = (b_{i+1}/b_{i-1})^{a_i} \mod p. \quad (3)$$

3: Cloud sends $c_1, \cdots, c_n$ to each $U_i, i = 1, \cdots, n$.

4: Each $U_i, i = 1, \cdots, n$ computes the random number

$$R_i = (b_{i-1})^{na_i} \cdot c_i^{n-1} \cdot c_{i+1}^{n-2} \cdots c_{i-2} \mod p. \quad (4)$$

Although each user generates R_i individually, all R_i equal to the same random number

$$R = g^{a_1a_2 + a_2a_3 + \cdots + a_na_1} \mod p. \quad (5)$$

Algorithm 2 Scramble code generation.

Input: Vector dimension N; Random number R; Sorted set $S = \{1, 2, \cdots, N\}$;

Output: Scrambled dimension sequence $\mathbf{SC}$;

1: **for** $k = N - 1$; $k >= 0$; $k--$ **do**
2: $i = R/k!$;
3: $\mathbf{SC}[N - k] = S[i]$;
4: Remove $S[i]$ from S;
5: $R = R \mod k!$;
6: **end for**
7: **return** $\mathbf{SC}$;

recalculation. When the membership changes, the cloud can insert/remove users into/from the exiting ring of Algorithm 1 to update the random number for a new round. Note that, in this case, most users do not need to recalculate the Step 2 in Algorithm 1. Based on our evaluation, the runtime for random generation time is usually only 0.014s, which is negligible for human movement. Once the random number is updated, the system achieves randomized transformation outputs for the same feature vector in different rounds.

VI. Prototype Implementation and Evaluation

A. Prototype Implementation

We implement prototype systems of both variants. To support automatic people detection, we implement the most popular face detection [13] and pedestrian detection [7] algorithm based on the library of OpenCV. To generate portrait graph, we adopt JSeg [14] for image segmentation. Leveraging the library of MPEG-7, a 48-byte eigenfaces vector [4] is extracted as the property of a node labeled with face, and a 64-byte color histogram vector and a 20-byte texture vector (edge histogram with 4 blocks and 5 orientations) are extracted as the property for other nodes. InvisibleMe is compatible with any other vector-based feature descriptors. Obtaining the matching results, invisible people are removed from the photo by blurring and inpainting [12], as shown in Fig. 4. Except the image processing, all other building blocks are realized using Java, including portrait graph matching, LSH, random number agreement, vector scrambling and the messaging module. The client side are developed as an app on Android system for case study. A user starts this app by inputting his/her portrait profile via selfieing.

B. Case Study and Experiment setup

To test the practicality and efficiency of InvisibleMe the evaluation is conducted in a crowded real-life scenario: a networking workshop with more than 50 attendees in a $200m^2$ meeting hall. 10 volunteers (4 female and 6 male) acted as invisible users and also photographers. Within one day, the volunteers took photos freely and our system recorded the cost and photos. After the experiment, we got 208 photos. 1326 pedestrians are detected which belong to 42 individuals (7 female and 35 male), but only 412 faces are detected. It implies that a whole body detection and description (*e.g.*, our graph model) is necessary. We manually labeled all captured people as the ground truth for the following evaluations.

Experiment setting. In the experiments, we use three types of phones as clients: HTC G10 (1024Hz CPU and 768M RAM), HTC G23 (1536Hz CPU and 1G RAM) and HTC New One (1741Hz CPU and 2G RAM). One laptop is used as the cloud: ThinkPad X1 with i7 2.7GHz CPU and 4GB RAM. Based on our extensive evaluation, to achieve the tradeoff between matching efficiency and accuracy, we set the parameter of the matching methods as $\xi_s = 0.5$ and the parameters of LSH as $W = 3$ and $m = 128$ (the hashed vector is 128 bit). For the random number generation, N is set to 512, which provides sufficient protection for the random number agreement protocol. The analysis of parameter setting is omitted due to space limitation. The the following evaluation, we denote the implementation of the baseline system as InvisibleMe-Basic and the implementation of the advance system as InvisibleMe-Advanced .

C. Matching Accuracy

Here we investigate the most important metric, the portrait matching accuracy of both variants, which determines the correctness of invisible people removal.

We start by examining the consistency and distinguishability of user's portrait graph by self-similarity (similarity between the same entity's portrait graphs) and cross-similarity (similarity between different entities' portrait graphs). In this evaluation, we remove the face property since it is highly distinctive but cannot always be obtained. Figure. 6 presents the evaluation results using the dataset. The upper blue line stands for mean self-similarity for each entity, and the lower red line is mean cross-similarity between this entity and all other entities. We notice that, generally portrait graph has a good consistency, *i.e.*, high self-similarity and small variance. And the obvious gap between self-similarity and cross-similarity shows a good distinguishability. In fact, in most cases, portrait graph can provide accurate matching without face features, which implies better privacy protection.

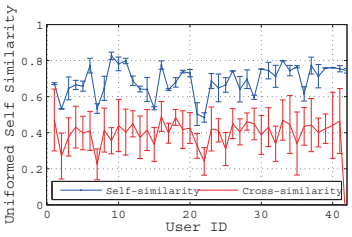

Fig. 6: Portrait similarity variances.

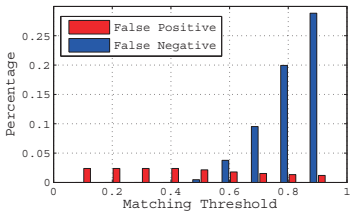

Fig. 7: FP and FN in basic scheme

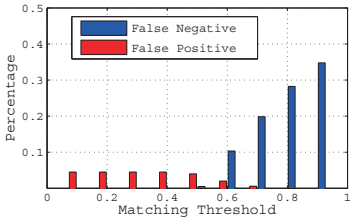

Fig. 8: FP and FN in advanced scheme

Then, we analyze the matching correctness by analyzing all possible combinations using the dataset. A false negative (FN) happens when a user A is invisible, but not removed from the photo due to a match score lower than a match threshold θ_s. A false positive (FP) happens when user A is invisible, but another visible user C is removed due to a higher match scores than both the threshold and A's score. In InvisibleMe-Basic matching is conducted on portrait graph with plain feature vectors. Fig. 7 illustrates the percentage of FN and FP changing with different threshold θ_s. By selecting the threshold $\theta_s = 0.5$, InvisibleMe-Basic achieves 0.5% false negative and 2.1% false positive without using any face property. With face property, the false negative decreases to about 0.1% and false positive is less than 1%. In InvisibleMe-Advanced feature vectors are transformed by scrambling and LSH. While the scrambling retains the accurate distance between vectors, LSH could cause some accuracy loss. Will the transformation reduce the matching accuracy? With appropriate parameters $m = 128$ and $W = 3$, InvisibleMe-Advanced achieves comparable accuracy with InvisibleMe-Basic , as shown in Fig. 8. When $\theta_s = 0.5$, the false negative is about 0.7% and the false positive is about 2.9% without any face property. So, both variants support accurate matching and our vector transformation achieves good portrait feature privacy protection with little accuracy loss. Based on the evaluation, in the rest experiments, the threshold θ_s is set to 0.5.

D. Micro Benchmark

Communication cost. In InvisibleMe-Basic each face node takes 48B and each other node takes 84B. The size of the portrait graph depends on the node number k. For most applications, $k \leq 10$ is sufficient, so the communication cost for each portrait is 0.82KB. In InvisibleMe-Advanced after encoding, each vector is hashed to 128 bits, which reduces the size of a portrait to 0.15KB. Protocol InvisibleMe-Advanced requires extra communication for random number agreement, which is only about 0.19KB. InvisibleMe costs each participant less than 1KB data transmission to enable portrait privacy protection. The cost for a photographer depends on the people number in the captured photo, but in most cases (with less than 10 people in photo), less than 10KB overhead is incurred, which is much less than a photo. In general, InvisibleMe achieves much smaller transmitted data size and better privacy protection than transmitting the image itself.

Computation cost. In InvisibleMe-Basic the computation cost is composed of portrait graph generation on the client side, and portrait matching on the cloud. InvisibleMe-Advanced costs extra computation for random number agreement and vector transformation by scrambling and LSH. The runtime is only about 3 ms to transform ten 64-dimension feature vectors. Table. I presents all the decomposed computing time. It shows that, the major computation delay is caused by image processing. For a participant, it only needs to be executed once for the profile setup; for a photographer, it needs to be executed for every captured photo. The runtime of portrait detection and segmentation depend on the resolution and complexity of the photo, but the detection and segmentation results are not sensitive to scaling. Hence, in our prototype all images are scaled to about 240,000 pixels. For the photographer, on average it takes about 0.4s to conduct face and pedestrian detection. Given a portrait image/subimage, the processing time of segmentation and feature extraction is about 2.6s. On average, there are 28.2 regions of each portrait.

TABLE I: Microbenchmarks of Runtime (in second)

Client			
Segmentation	0.5	2.4	8.1
Extraction	0.02	0.25	1.3
Random-Gen	0.012	0.014	0.017
Cloud			
Matching (basic)	0.015	0.037	0.079
Matching (advanced)	0.006	0.01	0.039
Random-Init	1.31	0.9	1.57

Compared with the image processing, the runtime of graph generation and matching is nearly negligible. On the client side, only extra 0.014s runtime is required for random number generation in InvisibleMe-Advanced . On the cloud side, the time needed to match a pair of portrait graph is only about 0.04s in InvisibleMe-Basic and decreases to 0.01s in InvisibleMe-Advanced due to the hashed feature vector. The cloud also needs 0.9s to generate system parameters for random number generation. The millisecond-level portrait graph transmission delay is negligible too. So the total computation delays for both variants are about 3s on the client and 1s on the cloud, which results a 4s system computation delay.

Now we've learned the magnitude of the time cost for each component, the overall delay also depends on the number of co-located invisible neighbors. With more active peers sending privacy requests, the matching cost will increase, but compared to the image processing cost, the matching cost on the cloud side is still quite small. Besides, the power consumption caused by our protocol (second-level computation) is much smaller than that caused by photo capturing itself.

E. Case Evaluation

We conduct the case based evaluation as described in Section VI-B. When there are invisible users in the photo, the false negative rate was about 1.4% and the false positive rate was 0.9%. But when there are no invisible users in the photo, the false positive rate raises to 4.9% due to the absent of any true match users, and the threshold 0.5 was not high enough to exclude all false match. And the average time for successful invisible people removal is about 4 seconds.

F. Compare with Alternative Solution

For comparison purpose, we also realize private Euclidean distance computation using a partial homomorphic encryption (Paillier encryption) in the SMC manner (e.g. the method used in [5]). Using the same computer and test images, the Paillier-based method takes about 0.5s for feature vector encryption and 1.8s for portrait matching between a pair of feature vectors. But with our approach, the transformation cost is negligible and the matching cost is only 0.01s. Besides, our method requires no interaction during the matching process. The comparison shows the a significant efficiency of our system.

VII. Related Work

Visual Privacy Protection There is a trivial solution to protect image content. Blacking out private contents, *e.g.* human faces, thwarts any possible violation of owners' privacy. For example, systems like Blinkering Surveillance [22] and [23] use computer vision methods to hide sensitive contents from video frame. But in a photographing scenario, the challenge is how to match people's privacy requests with people in the photo. Face recognition is a alternative way to solve the matching problem, *e.g.* Eigenfaces [4]. To use face recognition approaches, it requires the people to face to cameras. Besides, during the information exchange process, face descriptors could be leaked to adversaries. There are some work providing privacy-preserving face recognition leverages homomorphic encryption, by which a client can privately search for a specific face image in the face image database of a server, *e.g.*, [6]. Those methods provide privacy protection to the requested images as well as the outcome of the matching algorithm, but the computation overhead is large and the result is not secure against the service provider.

Graph matching. Graph matching methods have been applied in many tasks such as face recognition [15], fingerprint identification [16] and others [17]. Their application-oriented specifications, *e.g.*, assumptions about node numbers, graph structure and pre-knowledge of correspondences, make them difficult to be applied in this work.

Privacy-Preserving Distance Computation. Euclidean distance can be computed privately among parties using secure multi-party computation (SMC) methods [5], [6] or garble circuit [6]. However, they usually require online interactions among data owners. Moreover, their large computation cost and ciphertext size make them unsuitable for mobile applications. [24] proposes an approach using Fourier-related transforms to hide accurate data values and to approximately preserve Euclidean distances among them. It works well for some data mining purpose on large datasets, but the transformation is public and deterministic and it cannot prevent malicious user from dictionary attack.

VIII. Conclusion

In this work, we present a new approach InvisibleMe to protect users' portrait privacy during photo taking and sharing. With our system, users that are unwilling to be photographed will be automatically erased from the pictures in a lightweight and privacy-preserving way. To achieve this goal, we propose the integrated system model, a graph matching scheme to locate people in pictures and a privacy-preserving vector distance computation method. We have fully implemented our protocol, and thoroughly evaluated our design.

References

[1] "Business insider, http://www.businessinsider.com/seattle-bar-bans-google-glass-2013-3."

[2] "http://www.franken.senate.gov/?p=press_release&id=2699."

[3] J. Luo, Y. Ma, E. Takikawa, S. Lao, M. Kawade, and B.-L. Lu, "Person-specific sift features for face recognition," in *ICASSP*. IEEE, 2007.

[4] M. Turk and A. Pentland, "Eigenfaces for recognition," *Journal of cognitive neuroscience*, vol. 3, no. 1, pp. 71–86, 1991.

[5] J. Katz, A. Sahai, and B. Waters, "Predicate encryption supporting disjunctions, polynomial equations, and inner products," in *EUROCRYPT*, 2008, pp. 146–162.

[6] A.-R. Sadeghi, T. Schneider, and I. Wehrenberg, "Efficient privacy-preserving face recognition," in *ICISC*, 2010, pp. 229–244.

[7] B. Leibe, E. Seemann, and B. Schiele, "Pedestrian detection in crowded scenes," in *CVPR*. IEEE, 2005.

[8] P. Simoens, Y. Xiao, P. Pillai, Z. Chen, K. Ha, and M. Satyanarayanan, "Scalable crowd-sourcing of video from mobile devices," in *Mobisys*. ACM, 2013.

[9] "Face reconstruction, www.cs.princeton.edu/ cdecoro/eigenfaces/."

[10] T. Xu and Y. Cai, "Feeling-based location privacy protection for location-based services," in *CCS*. ACM, 2009.

[11] N. Komodakis, "Image completion using global optimization," in *CVPR*. IEEE, 2006.

[12] A. Criminisi, P. Pérez, and K. Toyama, "Region filling and object removal by exemplar-based image inpainting," *IEEE Transactions on Image Processing*, vol. 13, no. 9, pp. 1200–1212, 2004.

[13] P. Viola and M. J. Jones, "Robust real-time face detection," *International journal of computer vision*, vol. 57, no. 2, pp. 137–154, 2004.

[14] Y. Deng and B. Manjunath, "Unsupervised segmentation of color-texture regions in images and video," *IEEE TPAMI*.

[15] L. Wiskott, J.-M. Fellous, N. Kuiger, and C. Von Der Malsburg, "Face recognition by elastic bunch graph matching," *IEEE TPAMI*, vol. 19, no. 7, pp. 775–779, 1997.

[16] D. Isenor and S. G. Zaky, "Fingerprint identification using graph matching," *Pattern Recognition*, vol. 19, no. 2, pp. 113–122, 1986.

[17] N. Hu, R. M. Rustamov, and L. Guibas, "Graph matching with anchor nodes: A learning approach," in *CVPR*. IEEE, 2013.

[18] H. W. Kuhn, "The hungarian method for the assignment problem," *Naval research logistics quarterly*, vol. 2, no. 1-2, pp. 83–97, 1955.

[19] M. Datar, N. Immorlica, P. Indyk, and V. S. Mirrokni, "Locality-sensitive hashing scheme based on p-stable distributions," in *Symposium on Computational Geometry*. ACM, 2004.

[20] M. Burmester and Y. Desmedt, "A secure and efficient conference key distribution system," in *EUROCRYPT'94*, pp. 275–286.

[21] P. P. Lee, J. C. Lui, and D. K. Yau, "Distributed collaborative key agreement and authentication protocols for dynamic peer groups," *TON*, vol. 14, no. 2, pp. 263–276, 2006.

[22] A. Senior, S. Pankanti, A. Hampapur, L. Brown, Y.-L. Tian, and A. Ekin, "Blinkering surveillance: Enabling video privacy through computer vision," in *Security & Privacy*. IEEE, 2005.

[23] W. Zhang, S.-C. S. Cheung, and M. Chen, "Hiding privacy information in video surveillance system." in *ICIP*, 2005.

[24] S. Mukherjee, Z. Chen, and A. Gangopadhyay, "A privacy-preserving technique for euclidean distance-based mining algorithms using fourier-related transforms," *The VLDB Journal*, 2006.